

Why Real-world Planning is Difficult

Steve Chien, Randall W. Hill, Jr., and Kristina Fayyad

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive, M/S 525-3660
Pasadena, CA 91109-8099
{chien@aig, hill@negev, kristina@negev}.jpl.nasa.gov

Abstract

While work on AI planning systems dates back to the early 70s, applications of AI planning systems are few and far between. This paper describes several major obstacles to the fielding of planning systems related to: plan representation and usage, operational contexts of planning systems, and knowledge acquisition and maintenance for planning knowledge. We use examples drawn from our experiences in applying planning and decision support technology to two procedural-reasoning application tasks at JPL: an automated image processing task and an antenna control task.

1. Introduction

Why have so few actual planning applications been fielded? In this paper we describe a number of issues hindering such efforts, derived from working in two applications involving procedural reasoning: automated image processing system (called MVP - for **Multimission** VICAR Planner) and a decision support system for antenna operations (called LMCOA - for Link Monitor and Control Operator Assistant). We categorize these issues into three general classes. The first set of issues relates to more expressive representations for planning knowledge (such as more expressive action and temporal representations). Within this issue, we particularly highlight the importance of representing and reasoning about plan quality, as the need for this capability has arisen in both the MVP and LMCOA applications. The second issue is the exact operational context of the system. Most problems cannot be fully automated, hence there needs to be a natural mode of interaction between the user and the system--a clear and convenient division of labor and control between the user and the system. This issue of mixed initiative systems is one ignored by many planning paradigms. Finally, a key

This paper describes work performed by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. Other past and present members of the MVP team are Helen Mortensen, Christine Ying, Shouyi Hsiao, Alex Gray, Joc Nieten, and Jean L  rre. Other past and present members of the LMCOA team include Crista Smyth, Trish Santos, Kathy Sturdevant, and Lorraine Lee.

factor in determining the feasibility of automating a planning application is the issue of cost of knowledge acquisition, verification, and maintenance. As both the MVP and LMCOA applications involve the use of knowledge representations somewhat tailored to the applications and involved considerable knowledge acquisition and maintenance efforts, we briefly discuss these issues.

The rest of this paper is organized as follows. Section 2 provides background information on the MVP and LMCOA applications. Section 3 describes representational difficulties encountered in the MVP and LMCOA applications. Section 4 describes the difficulties of integrating a system into the MVP and LMCOA operational contexts. Section 5 outlines some of the issues relating to knowledge acquisition, knowledge verification, and knowledge maintenance relevant to the MVP and LMCOA applications.

2. The MVP and LMCOA Applications

We begin by providing an overview of the two applications which we use to illustrate our points. We first briefly describe the **Multimission** VICAR Planner application, in which planning techniques are used to automatically generate image processing programs from user specified image processing goals. We then briefly describe the LMCOA application, in which an automated reasoning system provides monitor, control, and decision support capabilities for operating Deep Space Network Antennas.

2.1 MVP: Automated VICAR Image Processing

The **Multimission** VICAR Planner (MVP)(Chien 1994a, 1994b) system is an Artificial Intelligence (AI) Planning system, which automatically constructs executable complex image processing procedures (using models of the smaller constituent image processing subprograms) in response to image processing requests made to the JPL **Multimission** Image Processing Laboratory (MIPL). The MVP system allows the user to specify the image processing requirements in terms of the various types of correction required. Given this information, MVP derives unspecified required processing steps and determines appropriate image processing programs and parameters to achieve the specified image processing goals. This

information is output as an executable image processing program which can then be executed to fill the processing request.

Currently, a group of human experts, called analysts, receive written requests from scientists for image data processed and formatted in a certain manner. These analysts then determine the relevant data and appropriate image processing steps required to produce the requested data and write an image processing program in a programming language called VICAR (for Video Image Communication and Retrieval!) (LaVoie et al, 1989). Unfortunately, this current mode of operations is extremely labor and knowledge intensive. This task is labor intensive in that constructing the image processing procedures is a complex, tedious process which can take up to several months of effort. This task is knowledge intensive in that it requires substantial knowledge of image processing, specifics of VICAR image processing programs, VICAR language constructs, and file and database organization and content. VICAR procedure generation is a common task - there are currently tens of analysts at MI PL alone whose primary task is to construct these VICAR programs. Many other users at JPL and other sites also write VICAR scripts, with the total user group numbering in the hundreds.

MVP2.0 is currently operational and in use by analysts at JPL's Multi mission Image Processing Laboratory (MIPL). MVP2.0 is written in C and operates with a Motif-based GUI on Sun workstations and an ascii-based interface on Vaxes under VMS. Over a test suite of 5 typical mosaicking and color reconstruction tasks, an expert analyst estimated that MVP reduces effort to generate an initial PDF for an expert analyst from 1/2 a day to 15 minutes, and that it would reduce the effort for a novice analyst from several days to 1 hour.

2.2 LMCOA and the Deep Space Network

The Jet Propulsion Laboratory manages a world-wide network of antennas, the Deep Space Network (DSN), which is responsible for providing the communications link with a spacecraft. Operations personnel are responsible for creating and maintaining the communications link by configuring the required subsystems and performing test and calibration procedures. This task of creating the communications link, known as **precalibration**, is a manual and time-consuming process which requires operator input of over a hundred control directives and operator monitoring of over a thousand event messages and several dozen displays to determine the execution status of the system. The existing Link Monitor and Control (LMC) system requires the operator to perform a large amount of textual keyboard entries, to monitor and interpret a large number of messages to determine the state of the system and to selectively cull out relevant information from dozens

.....
¹This name is somewhat misleading as VICAR is used to process considerable non-video image data such as MAGELLAN synthetic aperture radar data.

of pre-defined, data-intensive displays. This results in an environment in which it is difficult to operate efficiently.

The goal of the Link Monitor and Control Operator Assistant (LMCOA) task is to demonstrate automated operations techniques which will improve operations efficiency and reduce precalibration time. The LMCOA is a knowledge-based prototype system which incorporates Artificial Intelligence (AI) technology to provide semi-automated monitor and control functions to support operating the DSN 70-Meter antenna at the Goldstone Deep Space Communications Complex (DSCC). Improved operations is achieved by using a flexible and powerful procedural representation, by reducing the amount of operator keyboard entries and by providing explicit closed loop communications and control through an expert system module. An operational version of the LMCOA is currently being tested at a DSN 34-meter antenna station at Goldstone, California. The current prototype reduces the amount of operator inputs under nominal conditions from about 700 to less than 10 required for a mission. The focus of the current work is on dealing with anomalous conditions and unexpected interactions with the environment by the operator.

3. Representation Issues

Many of the obstacles hindering application of planning techniques to real-world problems can be characterized as representational difficulties. Of the representational issues we have encountered in the LMCOA and MVP applications, several can be attributed to the general area of representing and reasoning about plan quality. Other representation issues include representing complex actions and action effects.

For example, in the MVP application, an important concern is output image quality. For a planning system to be able to represent large portions of an analysts' expertise, the planner must be able to represent and reason about the effect of various image transformations on image quality. For example, one of the most commonly occurring image processing requests is for mosaicking, which is the process of combining a number of smaller images into a larger image (mosaicking). A frequent situation in mosaicking is that some of the images can be navigated absolutely - that is to say that some images contain features so that it is possible to determine exactly where each point on the image should be on the output image. However, the remainder of the images can only be correctly placed on the output image by matching up points which are common between them and other images (tiepoints).

For example, in planetary imaging applications, one might have 24 images taken of the planet Earth. Of these, 10 might have the edge of the planet visible in the image. For these 10 images it is possible to match the edge of the planet onto a geometric model of the planet to determine where the planet center is with respect to the image, thus accurately determining the final position of each of these 10 images on the final output image. However, for the remaining 14 images, they can only be placed by looking

for common geographic or atmospheric features between adjacent images in the overlap between images.

While the quality of each match can only be evaluated at runtime (by using pattern matching algorithms to find common features), the quality of these matches can be estimated by using the confidence in the match returned by the pattern recognition algorithm as well as comparing against the predicted overlay pattern computed from the spacecraft navigation and pointing information.

Ideally, an expert image processing planning system would be able to reason about these measures of image quality, to determine at runtime the best order in which to combine the images. Thus, by representing and reasoning about these combination operations and estimates of their impact on image quality we hope that MVP will be able to produce expert quality mosaics.

A secondary, but also important concern for MVP is the computational efficiency of the produced plan. If the image quality will be equivalent, there are sometimes different methods of achieving the same image processing goals but with different characteristics of computer runtime or disk storage. If MVP can reason about these types of costs for plans it will be able to produce plans which are more acceptable to the analysts and scientists.

Representing and reasoning about plan quality is also a key concern in the LMCOA application domain in at least two ways. First, overall execution time to setup (pre-calibration) and reset (post-calibration) the communications link subsystems should be minimized as this allows more data to be returned per operating time for the communications link. For instance, it can take up to two hours to manually pre-calibrate a DSN 70-meter antenna communications link for certain types of missions. Using the LMCOA, this time can be reduced to approximately thirty minutes, where further reductions in set-up time are limited by physical constraints of the communications link subsystems themselves. This reduction in operations time can save thousands of dollars each time the precalibration is performed, and for this reason efficiency is a primary measure of plan quality. Much of the efficiency of a plan is achieved by exploiting parallel plan path execution where the control of multiple subsystems is involved.

Another measure of plan quality is its robustness, that is, its ability to be used successfully in a broad range of situations. There are several aspects of plan robustness worth noting: generality, flexibility, and expressiveness. Before discussing each of these aspects of plan robustness, an overview of the representation being used in the LMCOA is first provided.

The LMCOA uses a temporal dependency network (TDN) to represent and automate LMC operations procedures. A TDN is a directed graph that incorporates temporal and behavioral knowledge and also provides optional and conditional paths through the network. The directed graph represents the steps required to perform an operation. Precedence relationships are specified by the nodes and arcs of the network. The behavioral knowledge identifies system-state dependencies in the form of pre- and

postconditions. Temporal knowledge consists of both absolute (e.g. Acquire the spacecraft at time 02:30:45) and relative (e.g. Perform step Y 5 minutes after step X) temporal constraints. Conditional branches in the network are performed only under certain conditions. Options paths are those which are not essential to the operation, but may, for example, provide a higher level of confidence in the data if performed. Each node in the TDN is called a "block" and contains actions to be performed. A block also has pre- and postcondition constraints and time tags associated with it. Ikn- (her details about TDNs are provided in (Fayyad & Cooper 1992).

Generality is the first aspect of plan robustness necessitated by the LMCOA domain. One of the missions frequently performed in the LMCOA domain is called the Ka-band Antenna Performance (KaAP) experiment. The KaAP TDN is currently implemented for the operational LMCOA; it is a generalized TDN in that it represents the many different ways that a KaAP experiment is to be executed. The support data for a particular KaAP experiment identifies a particular path through the TDN. For example, there is a data capture loop in the KaAP TDN which allows data to be captured from either a star or a planet, thus requiring different antenna modes. One experiment may specify acquiring data from the following sources in sequence: star 1, star2, star3. Whereas another experiment may specify acquiring data from: star], planet 1, star 1, star2. In terms of the cost of generating, maintaining and refining TDNs, a single generalized TDN is cheaper than several (really hundreds or thousands depending, on the number of different experiments) experiment-specific TDNs.

Flexibility is another aspect of plan robustness that has been a requirement in the LMCOA. For instance, the support data for a particular experiment may specify a particular path through the TDN, however, the operator also has the flexibility to alter this path in real-time. The TDN and LMCOA must be able to handle these real-time changes. Some of the changes that the operator can make to the TDN are skipping blocks, deleting commands in blocks, adding commands in blocks, and editing time tags on blocks. This flexibility of the TDN affects the evaluation of pre and postcondition constraints on blocks. Depending on the modifications made to the TDN, it is very likely that preconditions, postconditions, and time tags will become invalid,

Finally, the plan representation must be expressive in order to provide robustness. The LMCOA application, the TDN representation was initially kept extremely simple, although it included parallelism. As the intricacies of a particular domain's procedure became evident, more expressive representations were required. These included loops, metric time, and actions with temporal scope, and their inclusion complicated the application. As a prototype, the LMCOA became very specific to a particular TDN. For example, a "loop until time" construct was required. In such a case, the actions in the loop would be executed until the pre-specified time occurred. At that time, execution of

the loop would continue until a pre-specified exit point had been reached, such that the exit time of the loop was actually after the time specified in the looping construct. The alternative of abruptly executing the loop at a particular time is not always acceptable.

As a result of these representational issues, the TDN model is currently being reviewed and revised. It will no doubt become more complex, but also more robust. In addition, it will result in a more simple implementation of the LMCOA which executes TDNs. The resulting LMCOA will then be independent of any particular TDN.

As another representational difficulty, in the MVP application, large numbers of operator effects (as many as 50-100) with possible inconsistent parameter settings complicate the application. This means that MVP is often searching in the space of operator effects to determine an appropriate plan rather than the more typical planning case of searching in the space of operators or operator orderings. This proliferation of operator effects and presence of complex interactions among operator effects has required adaptation of traditional planning operator representations to include constraints among effects in operator effect preconditions. Thus, certain settings of program options will lead to certain operator effects - with there commonly being multiple program option settings to achieve the same effect. However, certain combinations of program options settings are incompatible. These incompatibilities are represented using bindings constraints on the program options. These explicit constraints allow MVP to perform least-cost search among program options (which control operator effects), resulting in an extremely efficient search among the possible program option settings. Thus, MVP is able to tractably search the large space of potential operator effects and program option settings.

4. Operational Contexts

Several of the difficult aspects of the MVP and LMCOA applications relate to what we call the operational context of the application system. In planning research, the planning problem typically is characterized as a batch problem, where the inputs and outputs of the system are carefully specified, and the planning system must produce a complete solution without user intervention. In the real world, this is rarely the case.

For example, in the MVP application, plans may be formed which require user inputs. For example, MVP may need to construct plans which involve determination of tiepoints between overlapping images (tiepoints are common reference points which appear in adjacent images and allow determination of points on one image relative to the other). In some cases the tiepoints can be determined automatically, but in other cases analyst intervention may be required to produce a high quality image. In the cases where human intervention is required, the planner must be able to reason about these cases, and be able to construct plans which are able to suspend execution as appropriate, waiting for analyst input and resuming when able.

Furthermore, occasionally, there are program parameters which may need to be adjusted by human analysts in a subjective fashion after inspecting the final image. In other more rare cases, the analysts may need to modify the produced image processing scripts to add further processing steps. Thus, because analysts must be able to modify MVP output, it is key that human analysts be able to understand and interpret MVP generated plans. In order to fulfill this requirement, MVP produces plans annotated with high-level comments, which detail at a conceptual level why MVP decomposed the problem in the manner it chose and which high-level goals are being attacked in which portion of the plan. This annotation greatly assists the analysts in understanding the structure of the produced image processing plans. At a lower level, the plan dependency structure itself can be used to explain the plan. This structure can be used to explain why certain image processing steps are needed, why certain parameters were set to the values used, or why image processing steps occur in the produced ordering.

The LMCOA application has to deal with several aspects of the operational context that affect planning: the domain is asynchronous, interactive, and real-time.

By asynchronous we mean that the effects of an action cannot be immediately observed and it may not have its intended effect. This affects the execution of the plan by forcing the LMCOA to monitor the state of the devices to which the plan's control actions have been sent. It must be able to recognize whether the action had its intended effect, and it must be able to deal with situations where the action had no effect at all or an unintended effect. For instance, an action may be sent to a device and there may be no response indicating that the action was received and executed. The LMCOA must take a corrective action once a time limit has passed for an effect to occur.

The LMCOA domain is interactive, meaning that the plan is not simply executed, rather, it is often necessary to re-plan or otherwise compensate for an interaction with the plan or the environment during its execution. External events, including human operator intervention, may interrupt a plan's execution. For example, if a problem occurs with a piece of equipment, the operator may need to take actions outside of what is represented in the TDN--especially in cases where the corrective action is beyond the control of the LMCOA. If the recovery actions take an extended amount of time, there may not be enough time to perform an planned equipment performance test as well as starting the acquisition of data at the required time. In this case a tradeoff must be evaluated. For example, should the data be captured without doing the performance test? Or would the data be useless without the performance test? Or is the data capture period long enough such that a certain amount of data can be lost at the expense of doing the performance test? Once the operator's intervention has been completed, the LMCOA may be given the command to continue to execute the TDN in the new context, which can be quite difficult since preconditions and postconditions

may have implicitly changed for the reasons above relating to time and priority.

The domain is real-time in that there are temporal constraints on the achievement of a plan's goals, which forces the LMCOA to continually monitor the plan's execution status as well as progress toward achieving the plan's goals. As previously indicated, the temporal constraints of the domain have to be taken into account when making decisions about re-planning after a plan failure.

S. Knowledge Acquisition and Knowledge Base Maintenance

One of the key elements in determining the feasibility of fielding a planning application is an assessment of the amount of effort required to construct the knowledge base and update and maintain the planning knowledge base. This has been particularly true in our experiences with the LMCOA and MVP applications. As a result, we have expended considerable effort in customizing the knowledge representations used for these applications and developing tools to facilitate knowledge base development and maintenance.

In the MVP application, we have developed two types of tools to assist in knowledge base development and maintenance. Static analysis tools analyze the knowledge base to detect simple cases where goals cannot be achieved. These cases are flagged and the user notified of these pathological cases. Completion analysis tools allow the user to detect cases where plans were almost able to be completed, but a certain subgoal could not be achieved or a certain protection could not be enforced. Completion analysis tools allow the user to quickly focus his attention on a specific portion of the knowledge base. These tools are described in further detail in (Chien 1994c).

In the LMCOA application, currently, the process of building the knowledge bases to represent a single TDN is manual and tedious. An LMCOA developer or knowledge engineer interviews operations personnel and scientists in order to obtain the details of a procedure. This information is then translated into a TDN, specified in terms of a graphical representation of the procedure, order of blocks, contents of blocks, precondition and postcondition constraints of blocks, as well as other support data. Much of this information must be specified in more than one place in the LMCOA, resulting in an overly complex implementation.

Several tools are under development to assist in the acquisition and maintenance of the plan knowledge base (for more details see Hill et al 1994). A TDN authoring tool is being developed to automate the specification of TDNs. Developers as well as operations personnel will be able to graphically specify the TDN and its contents. TDNs can be composed from parts of existing TDNs and libraries of actions at the least. A database will efficiently store a complete specification of a TDN as part of a TDN library. The same database will serve as a central repository for the TDN in the LMCOA, thus simplifying the LMCOA

implementation. In addition, the TDN authoring tool will include the capability to verify certain aspects of the TDN such as incompatible block ordering based on pre and postcondition constraints of blocks. The knowledge engineering effort for the LMCOA prototype is described in more detail in (Fayyad, Hill, and Wyatt, 1993).

Besides the TDN authoring tool, two other tools, RIDES and REBUS, are being developed and used for knowledge acquisition. The RIDES simulation authoring toolkit (Munro et al., 1993) is used to capture device models of the communications link equipment and subsystems. Besides using these models in the planner, the simulator also permits us to test the LMCOA's ability to cope with the operational context issues described in the previous section. REBUS, which stands for Requirements Envisioning By Utilizing Scenarios, (Zorman, 1994) is used to capture knowledge about the domain by using different scenarios to provide contextual information needed for planning. This provides us a way to understand how the subsystems controlled by the LMCOA actually work and how to control them under both normal and anomalous conditions.

Development of tools to facilitate in knowledge acquisition, verification, and maintenance are of prime importance. Furthermore, enriching the basic planning knowledge with sufficient knowledge to allow explanation in the course of training users and typical use is of key importance. These requirements are likely to require algorithms and techniques specialized to the particular representations used by MVP and LMCOA.

6. Summary

In summary, we have briefly described a number of issues which have complicated application of planning technology to two planning related projects at JPL: MVP - a planning system for automated generation of image processing procedures; and LMCOA - an intelligent system for assistance in antenna operations. These issues include expressiveness of representations (in particular to represent plan quality), operational contexts, and knowledge acquisition, verification, and maintenance.

References

- (Chien, 1994a) Chien, S. A., "Using AI Planning Techniques to Automatically Generate image Processing Procedures: A Preliminary Report," Proceedings of the Second International Conference on AI Planning Systems, Chicago, IL, June 1994.
- (Chien, 1994b) S. A. Chien, "Automated Synthesis of Image Processing Procedures for a Large-scale Image Database," Proceedings of the First IEEE International Conference on Image Processing, Austin, TX, November 1994.
- (Chien, 1994c) S. A. Chien, "Towards an Intelligent Planning Knowledge Base Development Environment," Proceedings of the 1994 AAAI Fall Symposium on

Planning and Learning: On to Real Applications, New Orleans, LA, 1994.

(Chien & Mortensen, 1993) Chien, S. A., and H. B. Mortensen, "MVP Domain Specification Document: Version 1.0," Unpublished Internal JPL Manuscript,

(Fayyad & Cooper, 1992) Fayyad, K. E. and L. P. Cooper. "Representing Operations Procedures Using Temporal Dependency Networks, " SpaceOPS '92, Pasadena, CA, November 1992.

(Fayyad, Hill, & Wyatt, 1993) Fayyad, K. E., R. W. Hill, Jr. and E. J. Wyatt, "Knowledge Engineering for Temporal Dependency Networks as Operations Procedures", AIAA '93, Computing in Aerospace 9 Conference, San Diego, CA, October 1993.

(Hill & Lee, 1992) Hill, R. W., Jr. and L. F. Lee, "Situation Management in the Link Monitor and Control Operator Assistant (LMCOA)", SpaceOPS 92: Second international Symposium on Ground Data Systems for Space Mission Operations, Pasadena, CA, 1992.

(Hill et al., 1994) Hill, R. W., Jr., Fayyad, K., Santos, T., Sturdevant, K. "Knowledge Acquisition and Reactive Planning for the Deep Space Network." Proceedings of the 1994 AAAI Fall Symposium on Planning and Learning: On to Real Applications, New Orleans, LA, 1994.

(LaVoie et al. 1989) S. LaVoie, D. Alexander, C. Avis, H. Mortensen, C. Stanley, and L. Wainio, VICAR User's Guide, Version 2, Internal JPL Document D-41 86, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 1989.

(Lee & Hill, 1992) Lee, L. F. and R. W. Hill, Jr. "Process Control and Recovery in the Link Monitor and Control Operator Assistant," Space Operations, Applications and Research Symposium, Johnson Space Center, Houston, TX, August, 1992.

(Lee & Cooper, 1993) Lee, L. F. and L. P. Cooper, "Operations Automation Using the Link Monitor and Control Operator Assistant", AIAA '93, Computing in Aerospace 9 Conference, San Diego, CA, October 1993.

(Munro et al., 1993) Munro, A., Johnson, M. C., Surmon, I. S., and Wogulis, J. L., "Attribute-Centered Simulation Authoring for Instruction. "Proceedings of AIED 93, World Conference on Artificial Intelligence in Education, Edinburgh, Scotland; 23-27 August 1993.

(Zorman, 1994) Zorman, Lorna. "Requirements Envisaging by Utilizing Scenarios." Ph.D. Diss., in preparation, University of Southern California, 1994.